

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

UTILITY PATENT APPLICATION FOR:  
**METHOD AND APPARATUS FOR EXCHANGING THE CONTENTS  
OF REGISTERS**

Inventors:

Kevin D. Safford  
4100 Suncrest Drive  
Fort Collins, CO 80525

Patrick Knebel  
2201 Greenmont Court  
Fort Collins, CO 80524

HP Docket No. 10992268

1      I. FIELD

3      The present invention relates to digital computer systems, and more particularly, but not  
by way of limitation, to methods and apparatus for executing instructions in such systems.

5      II. BACKGROUND

7      In x86 computer systems, the floating point unit (FPU) comprises a plurality of data  
registers. Floating point instructions treat this plurality of data registers as a register stack. All  
9      addressing of the data registers is relative to the register on the top of the stack. The register  
number of the current top-of-stack register is stored in a stack TOP field. Thus, load operations  
decrement TOP by one and load a value into the new top-of-stack register, while store operations  
11     store the value from the current top-of-stack register in memory and then increment TOP by one.

13     Many floating point instructions, however, only operate on the top one or two registers of  
a register stack. Thus, if the desired information is located in, e.g., the fourth stack register, one  
15     or more operations must be performed before the information in the fourth stack register can be  
moved into the top register of the stack where it can be operated upon. This creates a  
17     “bottleneck” in the stack. To this end, the floating point exchange register contents instruction  
19     (FXCH) is used in the IA-32 computer architecture to exchange the floating point information in  
a selected stack register with that in the top register of the stack. For example, the instruction

21     FXCH ST(0), ST(i)

will exchange the information in the top register in the stack (denoted ST(0)) with the *i*th register  
in the stack (denoted ST(*i*)). In this way, the bottleneck in the register stack can be alleviated by  
putting desired information at the top of the stack, where it can then be operated upon by most

1 floating point instructions. More information regarding the FXCH instruction may be found in  
the *Intel Architecture Software Developer's Manual, Volumes 1-3*, which are hereby incorporated  
3 by reference.

In many computer architectures, instructions, such as the FXCH, must be executed by  
5 emulation because the native hardware that supports such an instruction is not present. One way  
of emulating the FXCH instruction in such architectures is through a technique called register  
7 renaming. In register renaming, the physical registers in question (e.g., ST(0) and ST(i)) are  
mapped into a stack register map. To exchange the contents of the two physical registers, the  
9 pointers that map the physical registers into the stack register map are changed or “re-pointed”  
from their original register to the other register, and thus the operation is performed. But, at least  
11 one problem with register renaming is that it requires that the pointers be stored in additional  
hardware which adds to the cost and complexity of the system as well as consuming valuable  
13 space.

Another way of emulating the FXCH is to sequentially execute at least three micro-code  
15 instructions as follows:

```
move temp := ST(0);  
move ST(0) := ST(i);  
move ST(i) := temp;
```

This is the traditional method of exchanging the contents of the register. This sequence of  
19 instructions uses a temporary register to switch the contents of the top register ST(0) and the *i*th  
21 register ST(i). This method of emulation, with its three micro-code instructions, consumes three  
register ST(i). This method of emulation, with its three micro-code instructions, consumes three  
23 times as many clock cycles as the single FXCH instruction and, in some cases, may consume  
even more, depending upon the latency associated with the move operations. Thus, there exists a

1 need for methods and apparatus for emulating the FXCH instruction without adding excess  
hardware and that consumes relatively few clock cycles. More generally, there exists a need for  
3 methods and apparatus for exchanging the contents of two registers in a relatively quick and  
efficient manner.

5

### III. SUMMARY

7 In one embodiment of the present invention, there is a processor based computer system  
having dependency checking logic and a register stack, wherein the system overrides the  
9 dependency logic such that move instructions associated with the stack registers may be executed  
in parallel. In another embodiment, the system operates such that it can be determined whether a  
11 stack underflow exception has occurred and if it has, the move instructions can be flushed, and a  
micro-code handler algorithm invoked that operates to allow execution of the move instructions  
13 in parallel without a stack underflow exception.

15 **IV. BRIEF DESCRIPTION OF THE DRAWINGS**

Figure 1 is a block diagram of a computer system including the present invention.

17 Figure 2 is block diagram of the processor of Figure 1.

Figure 3 is an illustration of pipelined or lockstep operations.

19 Figures 4 is a flow chart of the operation of the stack underflow fault micro-code handler  
of the present invention.

21

1      **V. DETAILED DESCRIPTION**

3      **A. Description of an Embodiment**

5      Figure 1 illustrates a computer system 10 in which the present invention may be  
7      implemented. The computer system 10 comprises at least one processor 20, main memory 30,  
9      and various interconnecting data, address, and control busses (numbered collectively as 40). An  
11     instruction set 50 (which may be a guest instruction set) and an operating system 60 may be  
13     stored in main memory 30. As illustrated in Figure 2, the processor 20 comprises a floating point  
unit 70, dispatch or dependency checking logic 80, at least two execution units 90a and 90b,  
micro-code ROM 100, a register stack 120 (in this embodiment the register stack 120 comprises  
eight individual registers 120(0)-(7)), a floating point tag word (FPTW) register 130, and various  
busses and interconnections (numbered collectively as 110). (One skilled in the relevant art will  
note that there need not be a separate floating point unit — the execution units are equally  
capable of executing floating point instructions).

15     Instructions are provided to the processor 20 from main memory 30. The instructions  
17     provided to the processor 20 are macro-code instructions that map to one or more micro-code  
instructions 140 stored in the micro-code ROM 100. The micro-code instructions can be directly  
executed by processor 20. Also stored in the micro-code ROM 100 are a set of micro-code  
handlers 150 that may be invoked to handle processor exceptions.

19     The floating point unit 70 accesses the register stack 120 to store and retrieve data in  
response to instructions. The FPTW register 130 is updated accordingly.

21     The processor 20 may have a pipelined architecture and may have allow for parallel  
processing of certain instructions. Dependency checking logic 80 operates to determine which

instructions can be operated in parallel, i.e., whether to issue two instructions or one instruction per cycle to the execution units 90.

3

## **B. Method of Operation**

## 1. Parallel Execution of Move Instructions

Assume that the processor 20 is presented with the following instructions:

7                    ST(i): = move ST(0);  
                      ST(0): = move ST(i);

In a sequential microprocessor, this sequence of code is not capable of exchanging the contents of the two registers (ST(0) and ST(i). In a sequential microprocessor, each instruction is executed independently and in serial fashion. Thus, in this sequence the first instruction will overwrite ST(i) before the second instruction reads ST(i), and the result placed in ST(0) will be ST(0) rather than ST(i). This is obviously an incorrect result. As noted, the traditional method for correctly performing the desired exchange would require an additional temporary location and an extra instruction.

However, in the present invention, the computer system 10 emulates the FXCH

instruction by forcing the two move instructions to be executed in parallel (illustrated

19 conceptually below);

$\text{ST}(i) := \text{move } \text{ST}(0);$        $\text{ST}(0) := \text{move } \text{ST}(i);$

Because these two instructions each modify a register used by the other instruction, hardware would normally inhibit them from being executed in parallel. Thus, in the present invention, a mechanism is provided that forces the hardware to override the dependency checking logic, and

1 forces the hardware to execute the two instructions in parallel. Both instructions are provided to  
2 respective execution units 90a and 90b that (1) read their respective operand at substantially the  
3 same time, (2) write their results at substantially the same time (at least one clock cycle after the  
4 read), and (3) finish executing their instructions at substantially the same time (i.e., they operate  
5 in lockstep). The latency associated with the respective execution units 90a and 90b is such that  
6 this can be accomplished. This lockstep execution is illustrated conceptually in Figure 3. In this  
7 manner, the computer system 10 can effectively emulate the FXCH instruction in substantially  
less time than it would take to execute the three move instructions of a sequential system.

9

## 2. Stack Underflow Exception

11 In modern superscalar microprocessors (ones with the ability to execute multiple  
instructions in parallel) exceptions are precise, meaning that when an operation faults, all of the  
13 "younger" operations in the pipeline must be flushed. However, in the present invention, a  
temporary register is not used to hold intermediate results of operations in the pipeline -- i.e., the  
15 two instructions are completed in lockstep. Thus, if either one of the two move instructions  
above causes an exception, both of the move instructions are flushed and re-executed in parallel  
17 to prevent data corruption. However, that is not conventional operation for many  
microprocessors. That type of operation will work properly if the "first" instruction causes the  
19 exception, but not if the "second" instruction does -- in that case, conventional operation dictates  
that the "first" instruction would not be flushed (see example below):

21 ST(i) := mov ST(0) "first" 23 "older"	ST(0) := mov ST(i) "second" "younger"
--	---

1 Thus, in the present invention, the computer system 10 flushes both operations when either of  
them causes a fault.

3 An exception that may occur when these move instructions are executed is called a "stack  
underflow" exception. A stack underflow exception occurs when an operation attempts to read  
5 the contents of an empty stack register 120(0)-(7). A floating point tag word stored in the FPTW  
register 130 indicates whether a stack register 120(0)-(7) is empty or not. A defined architectural  
7 response to a stack underflow is to replace the empty register with a QNaN, mark it as non-  
empty, and then perform the instruction again. While it is possible to add hardware inside the  
9 execution units 90a and 90b to indicate which of the two move instruction caused the stack  
underflow fault, that is not desired because of the additional complexity and cost.

11 Thus, in the computer system 10, a micro-code handler algorithm 200, such as that  
illustrated in Figure 4, is invoked by processor 20 when a stack underflow exception occurs. At  
13 block 210 of Figure 4, a stack underflow exception has occurred when an attempt to execute the  
two move operations in parallel was attempted. At block 220, the micro-code handler algorithm  
15 200 causes the FPTW bits in the FPTW register 130 that correspond to the ST(0) register to be  
checked and then at block 230, a decision is made as to whether register ST(0) is empty or not. If  
17 register ST(0) is empty, at block 240 its contents are replaced with a QNaN and the  
corresponding FPTW bit is set to indicate that the ST(0) register is no longer empty. Proceeding  
19 now to block 260, emulation of the FXCH instruction is performed again (by issuing the two  
instructions in parallel again). If the register ST(0) was not empty, the exception must have  
21 occurred because register ST(i) was empty and, at block 250 the register ST(i) contents are  
replaced with a QNaN, the corresponding FPTW bit is set, and the emulation is performed again

1 at block 260. At block 270, if a stack underflow exception occurs once again, it is known that  
2 both registers involved in the operation must have been empty originally and that this time, the  
3 ST(i) register caused the exception. The registers contents at this stage are:

$$\text{ST}(0) := \text{QNaN} \quad \text{ST}(i) := \text{empty}$$

5 Accordingly, at block 280, the ST(i) register is loaded with a QNaN and at block 290, the  
6 emulation proceeds again, this time without any exceptions -- the QNaNs in the two registers will  
7 be harmlessly exchanged.

9           **C. Remarks**

10 By "forcing" the two move instructions to execute in parallel, emulation of the FXCH  
11 instruction may be achieved in substantially less time than previous methods and without adding  
12 hardware to the computer system 10. Furthermore, a microcode handler ensures correct  
13 execution of the emulated FXCH in the event of a stack underflow exception.

14 It will be readily apparent to those skilled in the art that innumerable variations,  
15 modifications, applications, and extensions of these embodiments and principles can be made  
16 without departing from the principles and spirit of the invention. For example, the methods and  
17 apparatus described herein may be applied to emulation of other instructions and to the handling  
18 of exceptions that occur when they are executed.

19           In another embodiment, rather than re-executing the instruction with respect to the ST(i)  
20 register (at block 260 of Figure 4), the computer system 10 instead checks the ST(i) register  
21 thereby eliminating the need to re-execute an instruction. In this embodiment, additional  
hardware or microcode is added to determine which register ST(i) is referenced in the instruction.

1 Accordingly, it is intended that the scope of the invention be only limited as necessitated by the  
accompanying claims.